ASCI Academia Strategic Alliances Program

Research Interests in Computer Systems and Computational and Computer Science Infrastructure

Presented at
Pre-Proposal Conference
December 5, 1996

By
Dona Crawford
Director, Distributed Information Systems
Sandia National Laboratories

ASCI success requires balanced growth of simulation components driven by the need for improved application performance

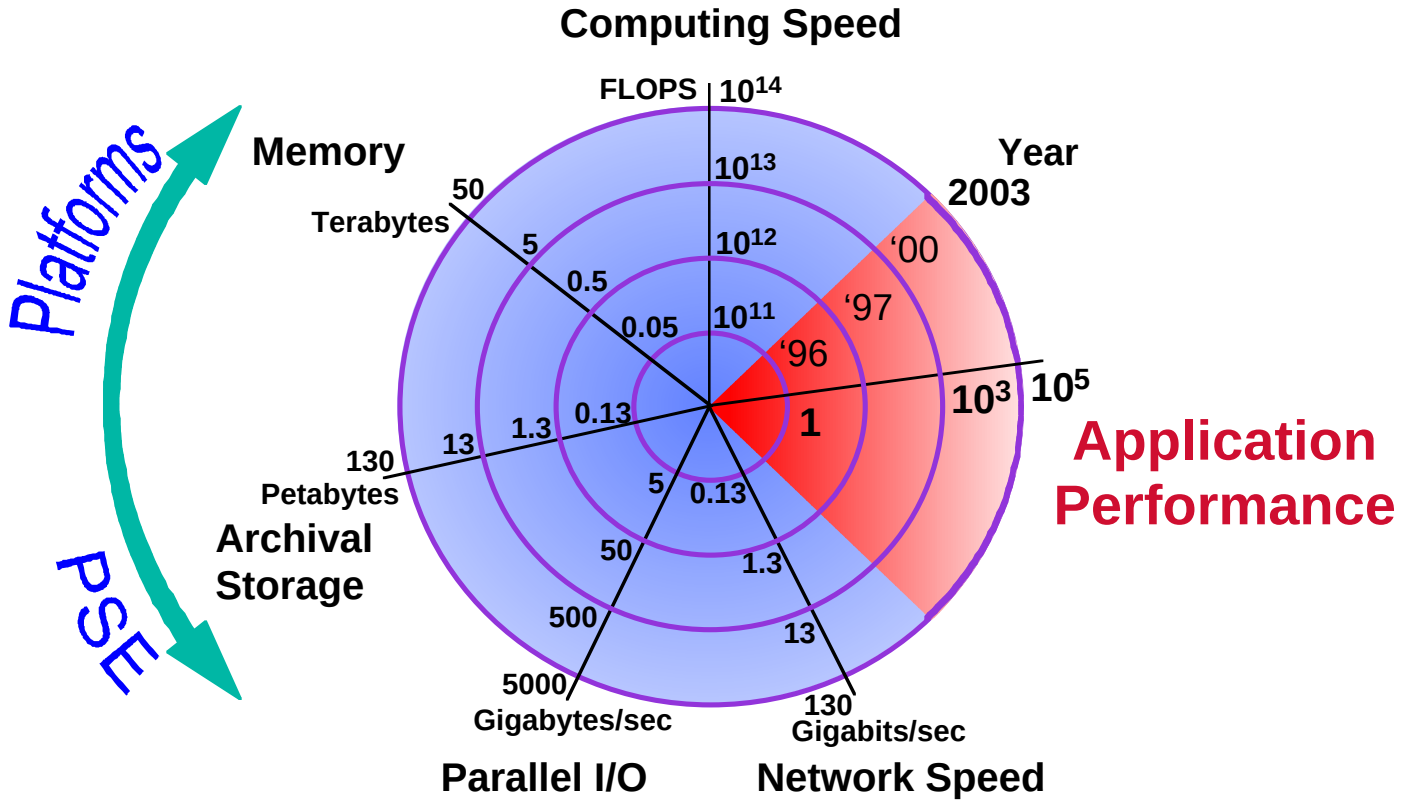

Our Applications Require Balanced Systems (Scale by 10^5)

For every 1 GigaFLOP peak performance, we need:

- **1 Gigabyte memory size**
- **50 Gigabyte disk storage**
- **10 Terabytes archival storage**
- **16 Gigabyte per second cache bandwidth**
- **3 Gigabyte per second memory bandwidth**
- **0.1 Gigabyte per second I/O bandwidth**
- **0.01 Gigabyte per second disk bandwidth.**
- **1 Megabyte/second archival storage bandwidth**

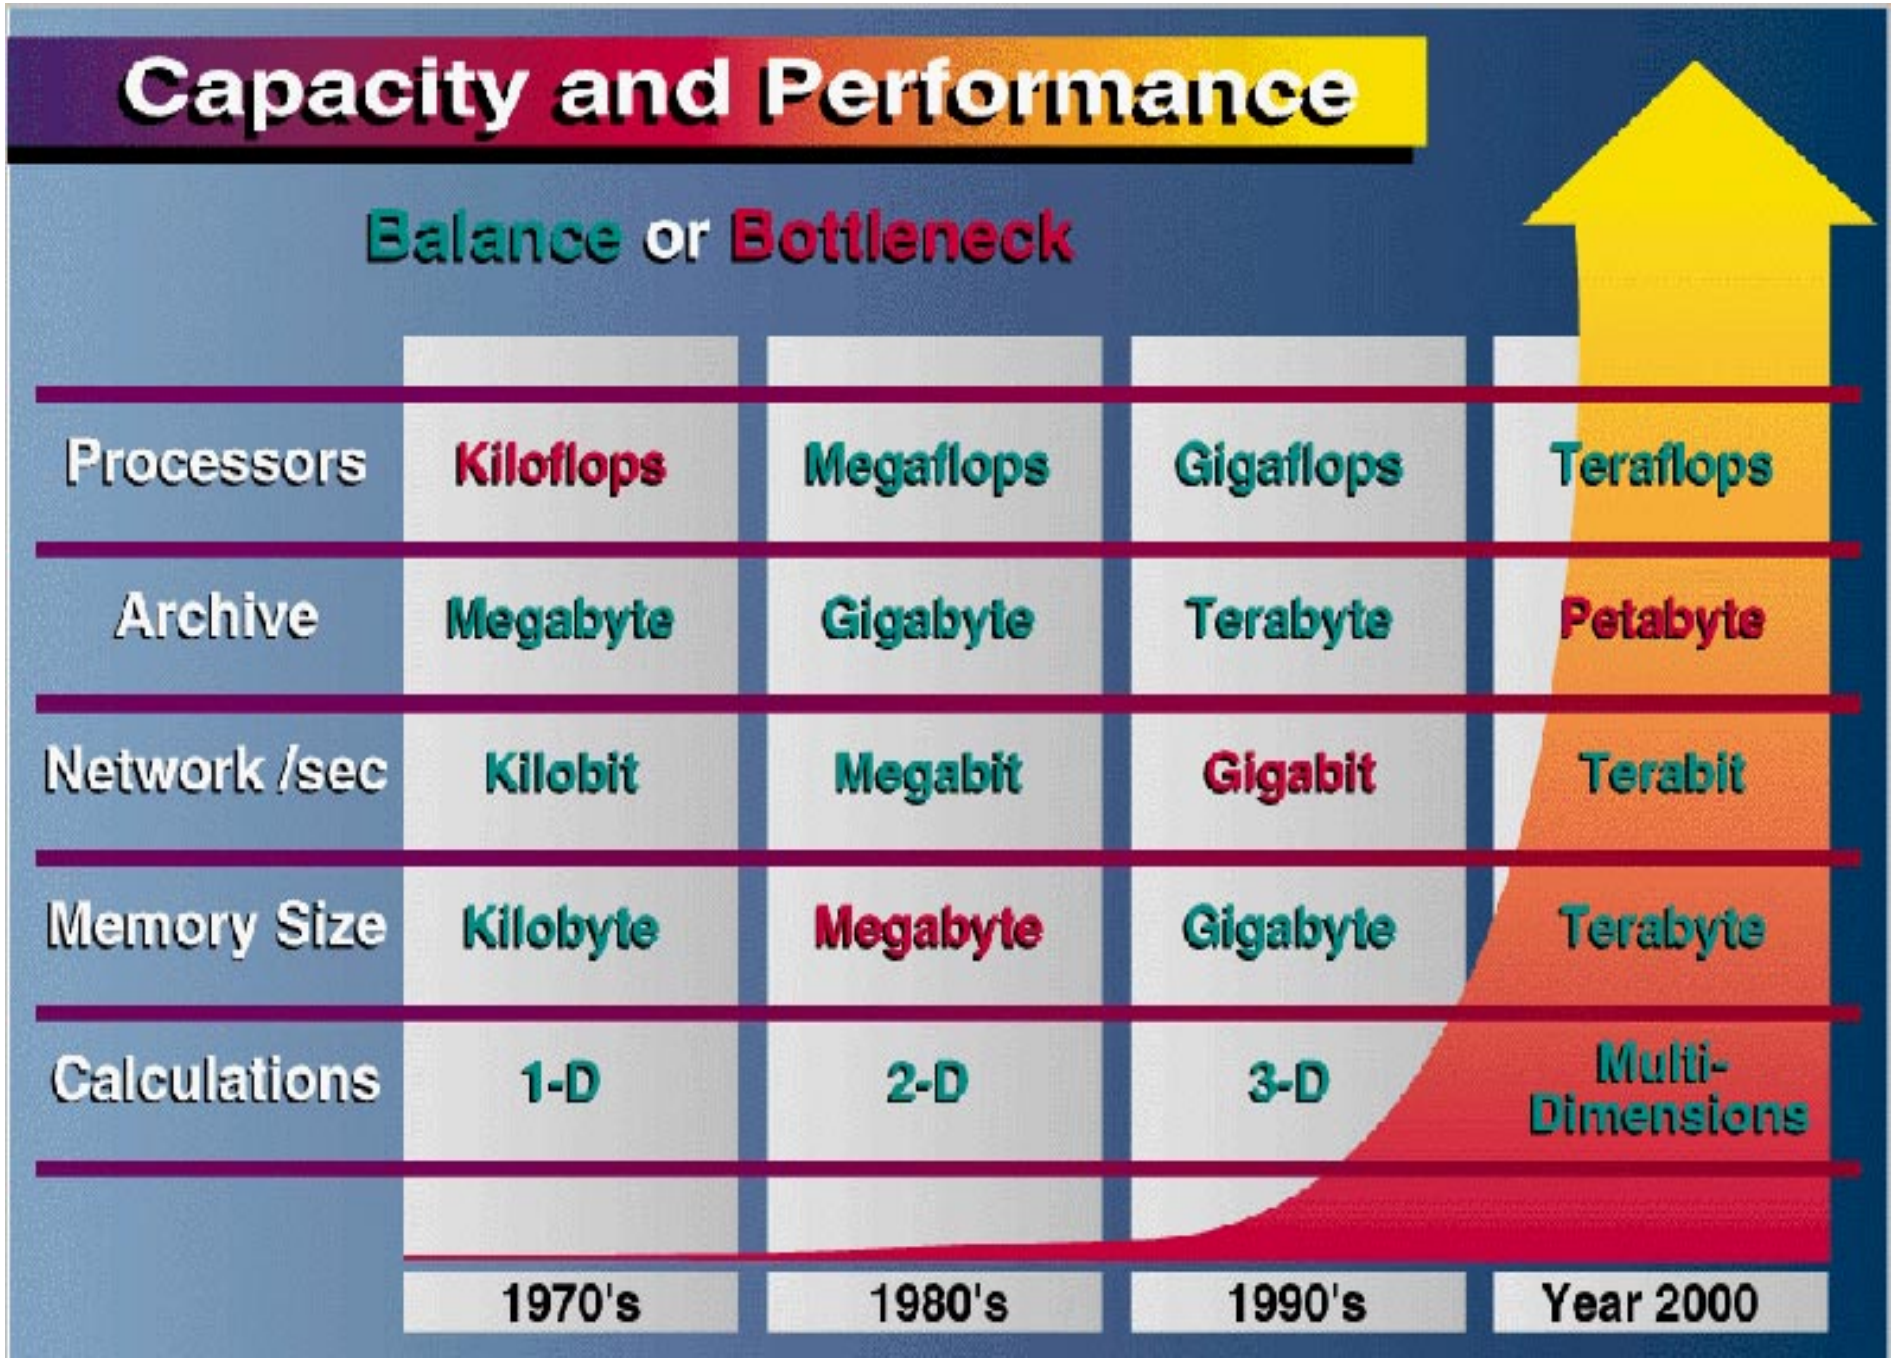
Capacity and Performance
Balance or Bottleneck
Processors
Archive
Network /sec
Memory Size
Calculations
Kiloflops
Megaflops
Gigaflops
Teraflops
Megabyte
Gigabyte
Terabyte
Petabyte
Kilobit
Megabit
Gigabit
Terabit
Kilobyte
Megabyte
Gigabyte
Terabyte
1-D
2-D
3-D
Multi-Dimensions
1970's
1980's
1990's
Year 2000

ASCI Hardware Requirements and Technology Trends

Level	Effective Latency (CPU cycles)	Bandwidth (Random read/write)	Size
On-chip cache**, L1	2-3 ●	16-32 B/cycle ●	10^{-4} B/flop * ● ↑
Off-chip cache**, L2 (SRAM)	5-6 ●	16 B/cycle ●	10^{-2} B/flop * ● ↑
Local main memory (DRAM)	30-80 (15-30) ↓	2-8 B/flop pk (2-8 B/flop sustained) ↓	1 B/flop ● ↑
"nearby nodes"	300-500 (30-50) ↓	1-8 B/flop (8 B/flop) ↓	1 B/flop ●
"far away nodes"	1000 (100-200) ↓	1 B/flop (1 B/flop) ↓	1 B/flop ●
I/O (memory disk)	10 ms ●	0.01-0.1 B/flop ●	10-100 B/flop ●
Archive (disk-tape)	Seconds ●	10^{-4} B/flop (0.001-0.01 B/flop) ↓	10^{2} B/flop 10^{4} B/flop ↓
User access	1/10 s (1/60 s)	OC3/desktop (OC12-48 /desktop) ↓	100 users ●
Multiple sites	1/10 s ●	●	●

Compute engine

Interconnect

* Equivalent integer and floating-point data calculation rates are required.

** Cacheless systems with equivalent performance are fully acceptable.

Primary investment priority

Secondary investment priority

1996-1998 Situation

(1998-2000 Requirements)

Industry Trend

↑ Industry gets better at meeting requirements

↓ Industry gets worse at meeting requirements

● Industry continues to meet requirements

ASCI Software Requirements and Technology Trends

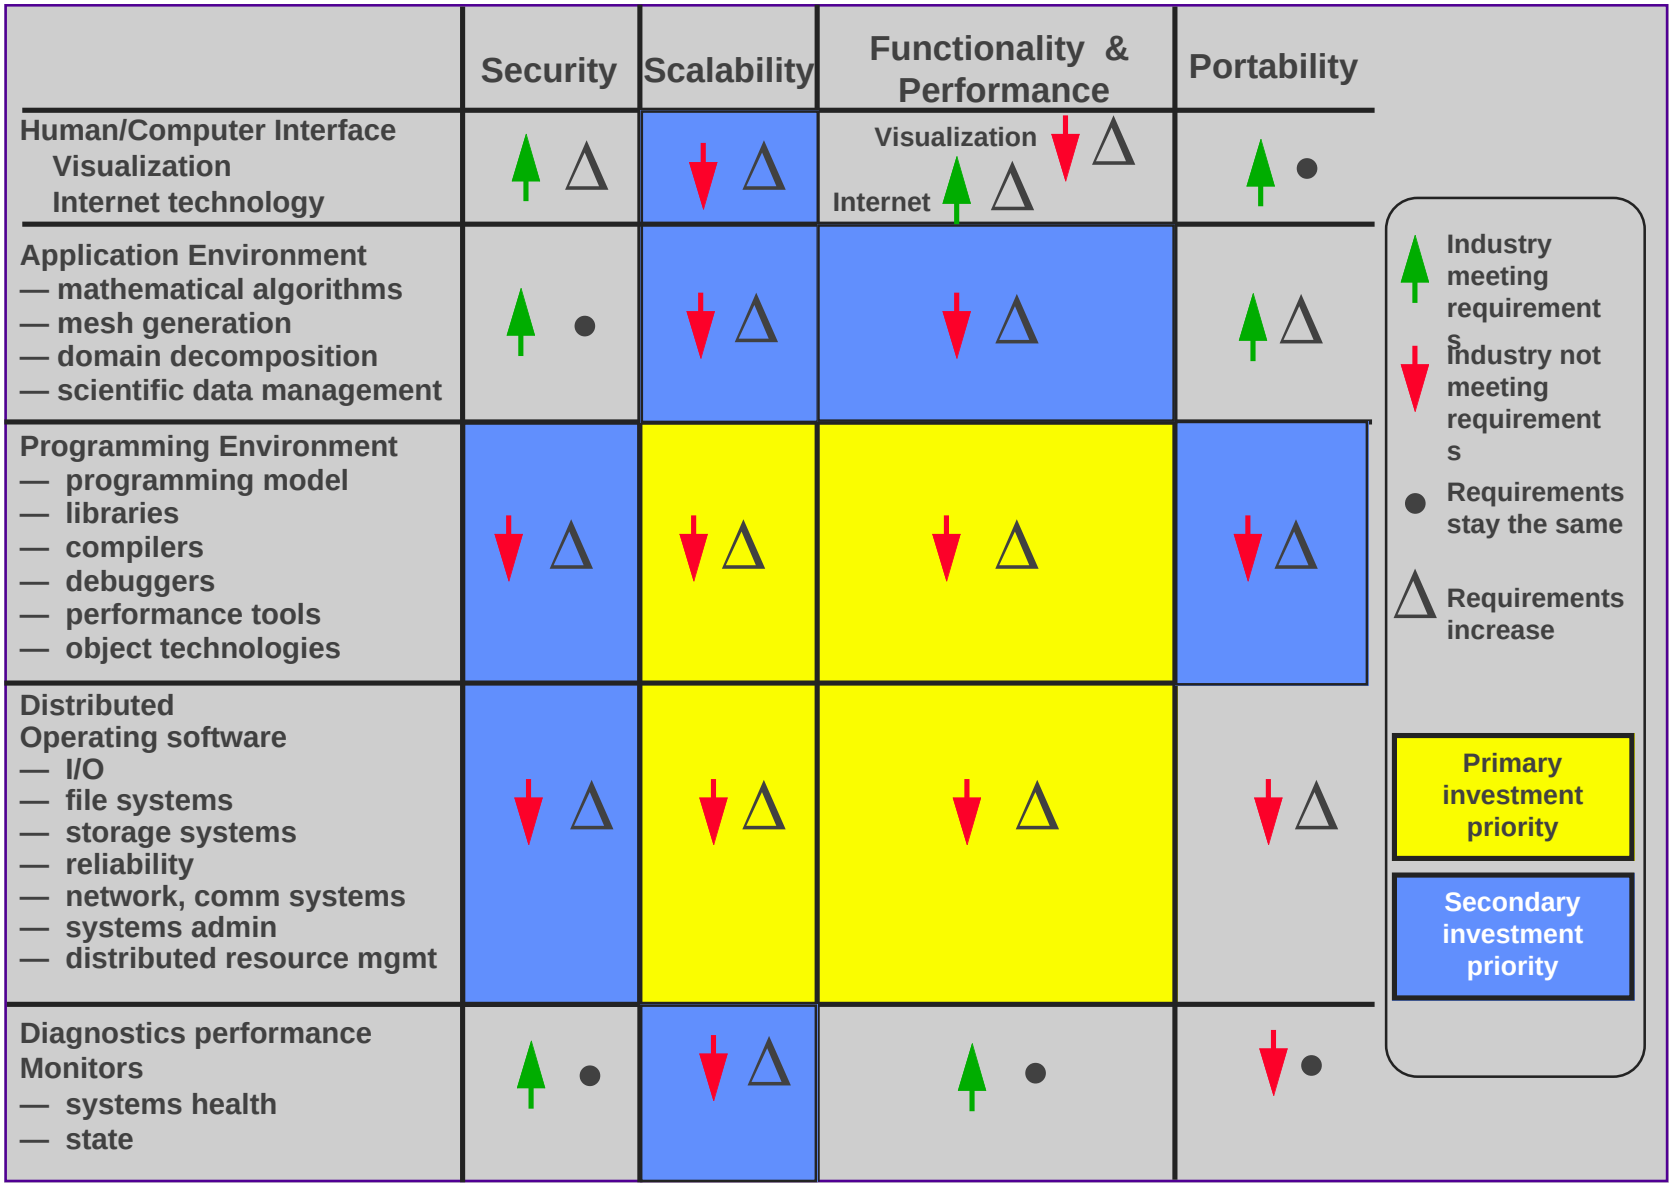

Example Hardware and Architecture Research Areas

- High Speed Interconnects
  - Order 10000 + cpu required
  - Congestion and bandwidth will likely cause current technology not to scale to this size
  - Create the illusion of shared memory
- WAN/LAN High Speed Networking
  - Gbs WANs/Tbs LANs
  - High Speed Encryption

Example Hardware Research Areas

- Storage Technology
  - Exabyte quantities of data will be stored
  - 100's of Gbs bandwidth to storage
    - Scalable network attached peripherals
    - RAIDs/RAITs
  - New Storage Technologies

Example Software Research Areas

- **Visualization**
  - **Scalable parallel visualization methods for terascale datasets**
  - **Hierarchical methods for representation and visualization techniques**
  - **Data-mining techniques for feature localization within the terascale data sets**
  - **Immersive visualization augmented by quantitative analysis capability**
  - **Integration with data archiving, retrieval, and analysis systems**

Example Software Research Areas

- **Computational Mathematics and Algorithms**
  - **Problem and Mesh Generation; Domain Decomposition**
  - **Scalable numerical methods and code frameworks are needed to enable terascale scientific simulations**
  - **Preconditioned iterative methods for sparse linear systems of equations**
  - **Methods for large systems of nonlinear equations**
  - **Methods for time-dependent differential equations, including ODEs, PDEs, and DAEs**
  - **Parallel adaptive mesh refinement libraries**
  - **Flexible frameworks for building codes**
  - **Application-aware communication routines and object classes**

New algorithms have yielded greater reductions in solution time than hardware improvements

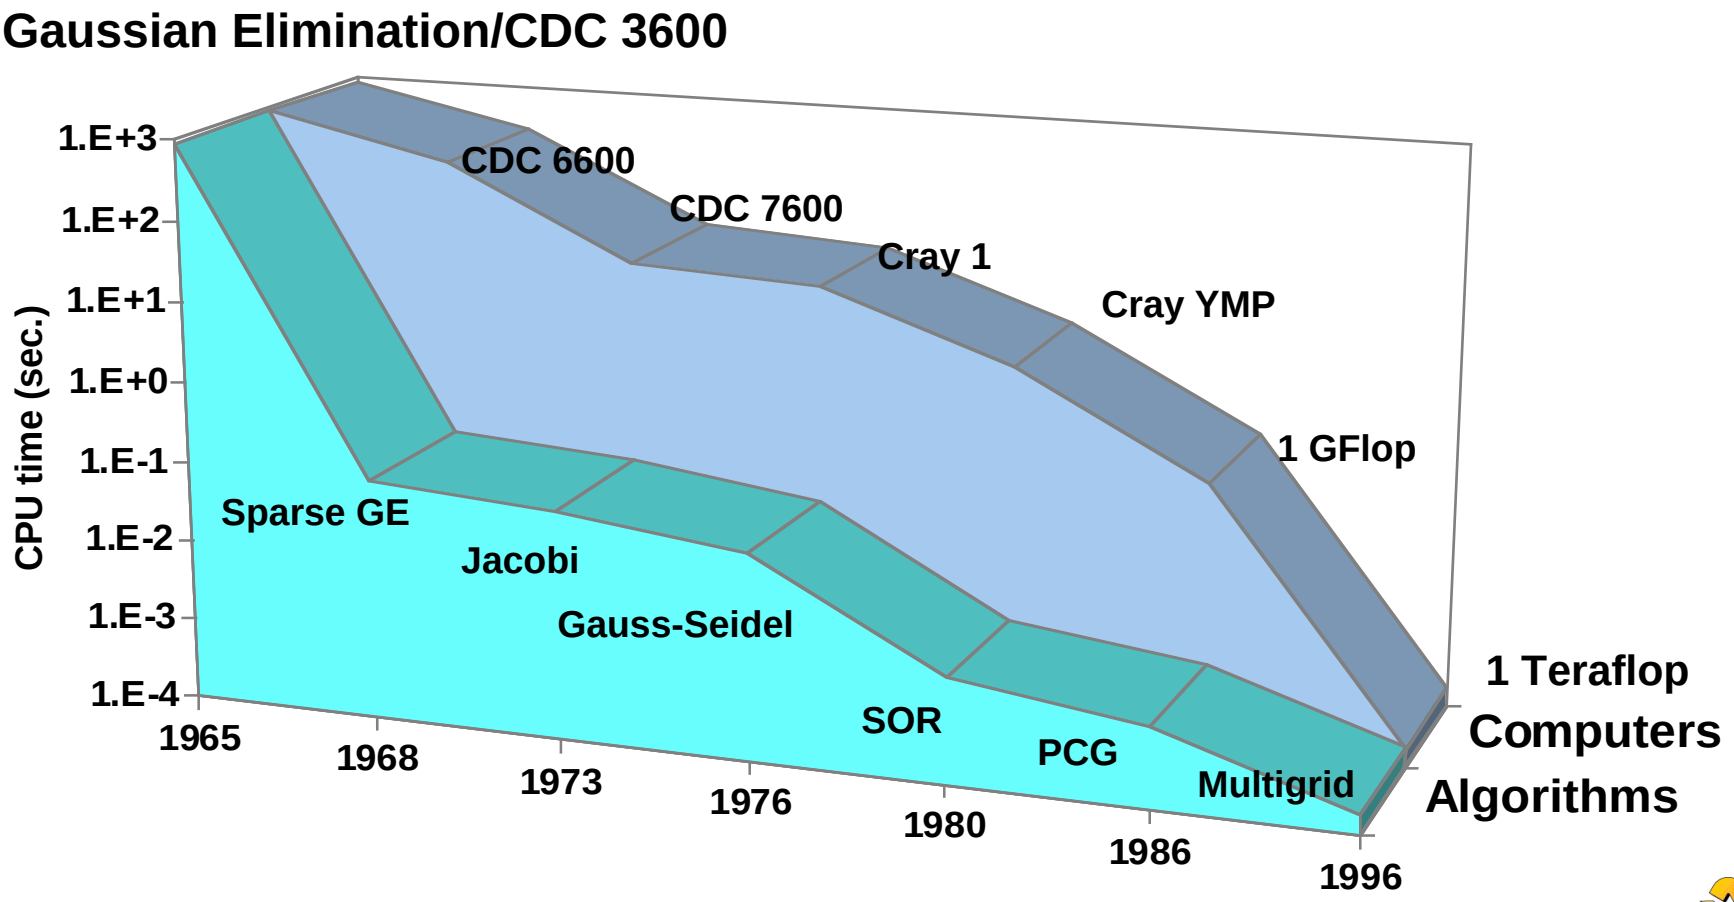

Example Software Research Areas

- **Programing Models/Compilers and Debuggers**
  - **Message passing (MPI) everywhere**
  - **Standard language extensions or compiler directives for shared memory parallelism, message passing for distributed,**
  - **Explicit threads for shared, message passing for distributed,**
  - **High level language constructs (HPF, HPC++, etc.)**
  - **Multilevel memory model**

High level programming models and abstractions are required that facilitate code reuse, reduce code complexity, and abstract away low level details necessary to achieve performance on a particular architecture.

Example Software Research Areas

- **Development Tools that scale to thousands of processes**
  - **Scalable parallel debuggers**
  - **Static analyzers**
  - **CASE tools**
  - **Verification and validation of the simulation results**
  - **Tools that work in a distributed environment**
  - **Tools to evaluate the performance and scalability**

Performance measurement, debugging, quality control, verification, and validation of codes becomes increasingly difficult as code size and complexity increases.

Example Software Research Areas

- **Distributed Scalable Operating Systems Software**
  - **Scalable, transparent parallel I/O (end-end: applications to tertiary storage)**
  - **Distributed file systems; archival storage systems**
  - **Operating systems will have to control single machine usage and parallel usage spread over many machines and across sites**
  - **Mechanisms for efficient thread creation, scheduling, and destruction (millions of threads)**
  - **Support the transparent creation, use, and maintenance of distributed applications in a heterogeneous computing environment**

Research is needed on both operating and programming system environments that provide services and tools to support the transparent creation, use, and maintenance of distributed applications in a heterogeneous computing environment.

Integration, Integration, Integration

- **It is expected that proposals will be developed that address the full range of tera-scale computing issues**

  (Physical Science, Computational Science, Computer Science)